Valid uncertainty quantification about the model in a linear regression setting

Ryan Martin,* Huiping Xu,† Zuoyi Zhang,‡ and Chuanhai Liu§

June 7, 2016

arXiv:1412.5139v2 [math.ST] 6 Jun 2016

Abstract

In scientific applications, there often are several competing models that could be fit to the observed data, so quantification of the model uncertainty is of fundamental importance. In this paper, we develop an inferential model (IM) approach for simultaneously valid probabilistic inference over a collection of assertions of interest without requiring any prior input. Our construction guarantees that the approach is optimal in the sense that it is the most efficient among those which are valid. Connections between the IM's simultaneous validity and post-selection inference are also made. We apply the general results to obtain valid uncertainty quantification about the set of predictor variables to be included in a linear regression model.

Keywords and phrases: Inferential model; post-selection inference; optimality; predictive random set; variable selection.

1 Introduction

1.1 Background

Linear regression is one of the most widely used statistical tools in scientific applications. Standard practice is to consider many predictor variables in hopes that only a few will identify themselves as being useful in explaining variation in the response variable. As a result, there is substantial uncertainty about the underlying model. The classical frequentist approach to deal with this problem is to use the data to help select a particular model, and there are a plethora of tools available, such as the Akaike information criterion (AIC, Akaike 1973), the Bayesian information criterion (BIC, Schwarz 1978), lasso (Tibshirani 1996, 2011) and its variants, including adaptive lasso (Zou 2006) and elastic net (Zou and Hastie 2005). Significance tests, such as those in Bühlmann (2013) and Lockhart et al. (2014), based on these or other methods have also been considered

*Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago, email: rgmartin@uic.edu

†Department of Biostatistics, The Richard M. Fairbanks School of Public Health and School of Medicine, Indiana University, email: huipxu@iu.edu

‡Regenstrief Institute, Inc., email: zyizhang@indiana.edu

§Department of Statistics, Purdue University, email: chuanhai@purdue.edu

recently. However, these frequentist methods provide no quantification of the uncertainty about the true model, i.e., they provide no way to conclude that one model being the true model is more plausible than another model being the true model.

Bayesian methods, on the other hand, are able to produce a summary of the uncertainty among the candidate models; see, for example, Clyde and George (2004) and O'Hara and Sillanpää (2009). Starting from a prior distribution on the set of possible models and a set of conditional priors on the model-specific parameters, a posterior distribution on the model space can be obtained via Markov chain Monte Carlo. Despite the conceptual simplicity of this approach, there are practical difficulties to overcome, including prior specification; in particular, real prior information is rarely available and improper default priors cannot be used. Without genuine prior information, the model uncertainty assessments coming from the corresponding posterior distribution are not guaranteed to be inferentially meaningful; see Section 1.2. Perhaps as a result of this lack of meaningfulness, the recent trend in the Bayesian literature on model selection is to de-emphasize uncertainty quantification, opting instead to report only a selected model. Therefore, modern Bayesian methods are effectively just frequentist selection procedures and they provide no summaries of model uncertainty.

Researchers have recently started to look beyond the classical frequentist and Bayesian schools for new solutions to challenging problems. See, for example, recent work on generalized fiducial inference (Hannig et al. 2015; Lai et al. 2015), confidence distributions (Xie and Singh 2013), and objective Bayes with default, reference, or data-dependent priors (Berger et al. 2009; Fraser 2011; Fraser et al. 2010; Martin et al. 2015). Our focus in this paper is the new perspective from the *inferential model* (IM) framework of Martin and Liu (2013). The key feature of this approach is that, for any assertion/hypothesis concerning the unknown parameter, it produces a probabilistic summary of the evidence in data for/against the truthfulness of that assertion. In addition to this internal or subjective interpretation, these inferential summaries are valid, i.e., suitably calibrated, facilitating an external or objective interpretation. The technical feature that distinguishes IMs from other approaches is the use of predictive random sets to produce a valid probabilistic summary of uncertainty about the parameter without a prior.

1.2 Lack of validity: an illustration

To further motivate our developments here, it will be helpful to elaborate on the claim above that the Bayesian approach does not, in general, provide an inferentially meaningful assessment of model uncertainty. We consider a simple illustrative example. Let $X = (X_1, X_2)$ be independent, with $X_i \sim \mathsf{N}(\theta_i, 1)$, $i = 1, 2$. There are four models in this case, one for each zero and non-zero combination for (θ_1, θ_2); here we will consider two "model assertions," namely, $\{\theta_1 = 0, \theta_2 = 0\}$ and $\{\theta_1 = 0\}$. A reasonable Bayesian model for this problem assigns a uniform prior over the four candidate models, and a $\mathsf{N}(0, v)$ prior for the model-specific parameters, where the to-be-specified variance, $v > 0$, controls the degree of prior uncertainty. The normal prior on the model-specific parameters is just a special case of the familiar g-prior in regression (e.g., Zellner 1986).

To quantify the uncertainty in the model given the observed data $X = x$, we report the posterior probabilities for the model assertions mentioned above. To be consistent with the notation and terminology in the rest of the paper, we will write these posterior

probabilities as $\mathsf{pl}_x(A)$, where A is one of the model assertions, and interpret this as the "plausibility" that model assertion A is true based on the observed data $x = (x_1, x_2)$. These posterior probabilities are proportional to sums of marginal likelihoods, i.e.,

$$\begin{aligned}\mathsf{pl}_x(\theta_1 = 0, \theta_2 = 0) &\propto \mathsf{N}(x_1 \mid 0, 1)\mathsf{N}(x_2 \mid 0, 1) \\ \mathsf{pl}_x(\theta_1 = 0) &\propto \mathsf{N}(x_1 \mid 0, 1)\mathsf{N}(x_2 \mid 0, 1 + v) + \mathsf{N}(x_1 \mid 0, 1)\mathsf{N}(x_2 \mid 0, 1),\end{aligned}$$

where the normalizing constant is the sum of all four marginal likelihoods. These are easy to compute, no sophisticated Monte Carlo methods are needed. The question is if the corresponding uncertainty quantification is meaningful. For example, the plausibility assigned to the true model should be large, but how large is large?

Figure 1 shows the distribution functions (gray) of these plausibility values for the two assertions of interest, for different true parameter values, and for three different values of v. It is natural to interpret probabilities on the $\mathsf{Unif}(0, 1)$ scale—e.g., 0.9 is a "large" value because it would be exceeded in only 10% of cases—so we take this as a reference. There are several key observations. First, there is a strong dependence on the (arbitrary) choice of prior variance; in fact, depending on v, even if the model assertion is true, as in Panels (a) and (c), it can happen that the posterior probability can never be "large." Second, there is no meaningful way that these distributions can be related to the $\mathsf{Unif}(0, 1)$ reference typically used for interpretation. Third, the way that shape of these distributions depends on the assertion and its truth/falsity is haphazard: the distributions in Panels (a) and (b) look indistinguishable even though the model assertions are true and false, respectively, and the distributions in Panel (c), based on a different but true assertion, have similar shapes; in Panel (d), the assertion is false, and the distributions look entirely different from the others. The take-away message is that the Bayesian posterior model probabilities need not have a meaningful interpretation in terms of quantifying uncertainty about the underlying model. The example here is relatively simple, so we can expect that similar issues will occur more generally.

As a preview, Figure 1 also shows the distribution function of the model plausibilities obtained from the optimal IM described in Section 4. The key observation is that the IM-based plausibility is valid in the sense that the comparisons with the $\mathsf{Unif}(0, 1)$ reference are consistent with our intuition in each case: when the asserted model is true (resp. false), the plausibility is stochastically no smaller (resp. no larger) than $\mathsf{Unif}(0, 1)$. This important difference compared to the Bayesian output is easiest to see in Panel (c). To our knowledge, only the IM approach proposed here provides valid model uncertainty quantification, so this is an important contribution.

1.3 Main contributions

Early work on IMs for regression is presented in Zhang et al. (2011). What is missing there is a rigorous treatment of model uncertainty quantification, and the present paper aims to fill this important gap. In particular, the main contributions here are two-fold.

First, we present new and general results on the construction of optimal predictive random sets for inference problems that involve multiple simultaneous assertions. Previous work on optimality (e.g., Martin and Liu 2013) focused on one assertion at a time but many important examples, such as the regression problem considered here, involve simultaneous consideration of multiple assertions. Our strategy is to decompose the set

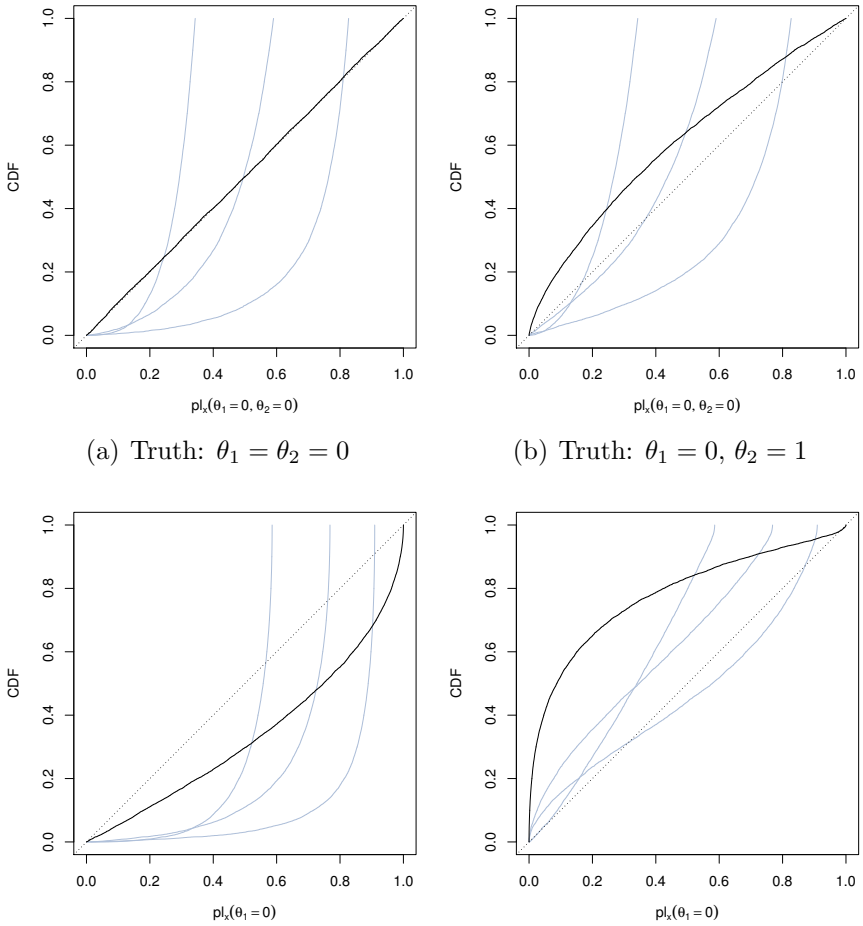

(a) Truth: $\theta_1 = \theta_2 = 0$

(b) Truth: $\theta_1 = 0$, $\theta_2 = 1$

(c) Truth: $\theta_1 = 0$, $\theta_2 = 1$

(d) Truth: $\theta_1 = 2$, $\theta_2 = 1$

Figure 1: Plots of the distribution function (CDF) of $\mathsf{pl}_X(\cdot)$ for different model assertions and several configurations of the true (θ_1, θ_2). Gray lines in each panel correspond to the Bayesian plausibility with prior variance $v = 1, 10, 100$; black line correspond to the IM plausibility described in Section 4; and the dotted line is the $\mathsf{Unif}(0, 1)$ reference.

of assertions under consideration into their basic building blocks, identify the optimal predictive random sets for these building block assertions, and then combine these individually optimal predictive random sets in a particular way to obtain the simultaneously optimal predictive random set for the collection of assertions.

Second, we apply these new results to construct an optimal IM for the important problem of model uncertainty quantification in regression. There we find that a hypercube predictive random set is optimal. The particular construction ensures that the IM's output is provably valid in the sense illustrated in Figure 1. To our knowledge, this is the first result on valid uncertainty quantification about the model in a regression context. As part of our development, we make an important connection between the IM's focus on simultaneous validity and modern attempts at post-selection inference, especially, that in Berk et al. (2013).

1.4 Outline of the paper

In Section 2 we review the basics of the IM approach in the context of the regression problem. Section 3 introduces the concept of efficiency and develops some general theory concerning the shape of the optimal predictive random sets for successively more general collections of assertions, concluding with a result (Theorem 4) that characterizes the IM for valid and optimal simultaneous inference over a class of so-called complex assertions. This result is specialized for model uncertainty quantification in regression in Section 4; there we draw connections between notion of simultaneous validity and post-selection inference, and show how the IM's validity property can be used to develop a variable selection procedure with provable control on certain frequentist error rates. The paper concludes with a short discussion in Section 5; proofs are deferred to the Appendix and some additional technical details are presented in the Supplementary Material.

2 IMs for regression

2.1 Baseline association

To set notation, consider the following specification of the linear regression model:

$$Y = X\beta + \sigma Z, \tag{1}$$

where $Y = (Y_1, \ldots, Y_n)^\top$ is a n-vector of response variables, X is an $n \times p$ matrix of predictor variables, $\beta = (\beta_1, \ldots, \beta_p)^\top$ is a p-vector of regression coefficients, σ^2 is the residual variance, and $Z = (Z_1, \ldots, Z_n)^\top$ is a n-vector of standard Gaussian errors, i.e., $Z \sim \mathsf{N}_n(0, I_n)$, a n-dimensional Gaussian distribution with mean zero and identity covariance matrix. We assume that $p < n$ and that X is fixed of full rank. Without loss of generality, we will assume that Y and the columns of X have been centered, which effectively marginalizes out the nuisance intercept parameter.

We refer to (1) as a baseline association between observable data Y (and X), unknown parameters (β, σ), and unobservable auxiliary variables Z. The importance of such an association can be seen as follows: if Z could be observed, then we could exactly solve (1) for (β, σ), leading to the "best possible inference." Since Z cannot be observed, a natural idea is to accurately predict the unobservable Z; see Section 2.4. It is clear that accurately predicting a high-dimensional quantity is more difficult than predicting a lower-dimensional quantity, so it is advantageous to reduce the dimension of the auxiliary variable as much as possible before carrying out the prediction step.

2.2 Auxiliary variable dimension reduction

Martin and Liu (2015a,b) discuss two distinct approaches for reducing the dimension of the auxiliary variable. The first is an approach based on conditioning, which is particularly useful in cases where the auxiliary variable is of higher dimension than the parameter. For example, in our regression problem, Z is n-dimensional while (β, σ) is $(p+1)$-dimensional, and $p + 1 \leq n$. Since the regression problem admits a $(p+1)$-dimensional minimal sufficient statistic, the IM dimension reduction is straightforward. Let $\hat{\beta}$ be the least-squares estimator of β, and $\hat{\sigma}$ the corresponding estimator of σ based

on the residual sum of squares. Then $(\hat{\beta}, \hat{\sigma})$ is a minimal sufficient statistic, and we can identify a lower-dimensional (conditional) association:

$$\hat{\beta} = \beta + \sigma V_1 \quad \text{and} \quad \hat{\sigma} = \sigma V_2 \tag{2}$$

where $V = (V_1, V_2)$ satisfies $V_1 \sim \mathsf{N}_p(0, M)$, $(n-p-1)V_2^2 \sim \mathsf{ChiSq}(n-p-1)$, independent, and $M = (X^\top X)^{-1}$. The key point is that we have replaced the n-dimensional auxiliary variable Z with a $(p+1)$-dimensional auxiliary variable V.

Here, as is typical in regression applications, σ^2 is a nuisance parameter. For such cases, it is possible to further reduce the dimension of the auxiliary variable via marginalization. Note that the association (2) can be rewritten as

$$\hat{\beta} = \beta + \hat{\sigma}(V_1/V_2) \quad \text{and} \quad \hat{\sigma} = \sigma V_2.$$

The general theory in Martin and Liu (2015b) says that the second equation above—the one that is free of β—can be ignored. This leaves a marginal association involving a p-dimensional auxiliary variable $W = V_1/V_2$. That is,

$$\hat{\beta} = \beta + \hat{\sigma} W, \qquad W \sim \mathsf{t}_p(0, M; n-p-1), \tag{3}$$

where the auxiliary variable distribution is a p-dimensional Student-t, with $n-p-1$ degrees of freedom, centered at the origin, with scale matrix $M = (X^\top X)^{-1}$. Again, the important point is that the $(p+1)$-dimensional auxiliary variable in (2) has been replaced by a p-dimensional auxiliary variable U. No further dimension reduction is possible.

It will be convenient to rewrite association (3) once more. Let D be a diagonal $p \times p$ matrix with the same diagonal as M. Then consider the association

$$\hat{\theta} = \theta + \hat{\sigma} U, \quad U \sim \mathsf{P}_U = \mathsf{t}_p(0, L, n-p-1), \tag{4}$$

where $\theta = D^{-1/2}\beta$, $\hat{\theta} = D^{-1/2}\hat{\beta}$, $U = D^{-1/2}W$, and $L = D^{-1/2}MD^{-1/2}$. Note, in particular, that $\theta_j = 0$ if and only if $\beta_j = 0$, so the question about which variables are included in the model has not been changed; also, L has all ones on the diagonal.

2.3 Predictive random sets

The interpretation of the association (4) from the standpoint of inference is that there is a particular value, say $u^\star$, for which that equation holds for the true θ. How we describe our uncertainty about this value $u^\star$ determines our uncertainty assessments about θ, given the observed data. What distinguishes the IM approach from Fisher's fiducial approach and its variants is the use of a random set, $\mathcal{S}$, to summarize uncertainty about $u^\star$. We refer to $\mathcal{S}$ as a *predictive random set*. The key to the IM approach's success is an appropriate choice of the predictive random set.

There is a rich mathematical theory of random sets, summarized nicely in Molchanov (2005). Here, for conceptual understanding, it will suffice to consider a discrete setting. Let $\mathbb{S}$ be a finite collection of subsets of a space $\mathbb{U}$, in our case, $\mathbb{R}^p$. This collection of sets will serve as the support for the random set $\mathcal{S}$; the individual sets in the collection $\mathbb{S}$ are called *focal elements*. Now, define a random set $\mathcal{S}$, supported on $\mathbb{S}$, simply by assigning probabilities to each of the focal elements; this characterizes the distribution, $\mathsf{P}_{\mathcal{S}}$, of $\mathcal{S}$. Of

course, defining a random set in a non-discrete case requires more care and, in particular, a notion of measurability of set-valued functions. The Supplementary Material provides some relevant technical background and explains how the random sets appearing in our IM development satisfy the required conditions automatically.

One is free to describe their uncertainty about the unobserved value $u^\star$ with any kind of predictive random set. However, the corresponding uncertainty assessments/inference about θ would be meaningful only if the two distributions in question, namely, P_U and $\mathsf{P}_{\mathcal{S}}$, have a suitable link. The following definition provides such a link.

Definition 1. A predictive random set $\mathcal{S}$, with distribution $\mathsf{P}_{\mathcal{S}}$, is called *admissible* if the following two conditions hold:

- the support $\mathbb{S}$, which is assumed to contain $\varnothing$ and $\mathbb{R}^p$, is *nested*, that is, for any two focal elements S and S' in $\mathbb{S}$, either $S \subseteq S'$ or $S' \subseteq S$, and
- the distribution $\mathsf{P}_{\mathcal{S}}$ of $\mathcal{S}$ satisfies

$$\mathsf{P}_{\mathcal{S}}\{\mathcal{S} \subset K\} = \sup_{S \in \mathbb{S}: S \subset K} \mathsf{P}_U(S), \quad K \subseteq \mathbb{U}. \tag{5}$$

Martin and Liu (2013) showed that, by choosing a predictive random set $\mathcal{S}$ that is admissible in the sense above, the corresponding IM output is valid, i.e., suitably calibrated for scientific inference; see Definition 2 below. Existence of a distribution $\mathsf{P}_{\mathcal{S}}$ that satisfies (5) is discussed in the Supplementary Material.

2.4 IM construction

First, the model provides a link between the observable data, the unknown parameters, and the unobservable auxiliary variables. Next, uncertainty about the unobserved value of the auxiliary variable, given the observed data, is encoded in a predictive random set. Finally, this uncertainty about the auxiliary variable gets passed through the association, yielding a corresponding uncertainty assessment about the parameter. For the dimension-reduced model in (4), the formal three-step IM construction is as follows.

A-step. Associate data $Y = y$ with unknown parameters θ and auxiliary variable U as in (3). This defines a set—in this case, a singleton set—of parameter values indexed by particular values of U, namely,

$$\Theta_y(u) = \{\theta : \hat{\theta} = \theta + \hat{\sigma} u\}, \quad u \in \mathbb{R}^p. \tag{6}$$

P-step. Predict the unobservable auxiliary variable U with an admissible predictive random set $\mathcal{S}$ with distribution $\mathsf{P}_{\mathcal{S}}$.

C-step. Combine the results of the A- and P-steps to get

$$\Theta_y(\mathcal{S}) = \bigcup_{u \in \mathcal{S}} \Theta_y(u).$$

Since $\mathcal{S}$ is a random set, so is $\Theta_y(\mathcal{S})$, and probabilities associated with this random set are used to summarize uncertainty about the unknown parameter. If $\mathcal{S}$ is such that $\Theta_y(\mathcal{S}) \neq \varnothing$ with $\mathsf{P}_{\mathcal{S}}$-probability 1 for all y—this is the only case we need to consider here,

but see Ermini Leaf and Liu (2012)—then, for a given assertion/hypothesis A about the unknown parameter θ, we calculate the belief function

$$\mathsf{bel}_y(A; \mathcal{S}) = \mathsf{P}_{\mathcal{S}}\{\Theta_y(\mathcal{S}) \subset A\}, \tag{7}$$

the $\mathsf{P}_{\mathcal{S}}$-probability that $\Theta_y(\mathcal{S})$ completely agrees with the assertion A. The other important quantity is the plausibility function

$$\mathsf{pl}_y(A; \mathcal{S}) = 1 - \mathsf{bel}_y(A^c; \mathcal{S}) = \mathsf{P}_{\mathcal{S}}\{\Theta_y(\mathcal{S}) \not\subset A^c\},$$

the $\mathsf{P}_{\mathcal{S}}$-probability that $\Theta_y(\mathcal{S})$ at least partially agrees with A. For given $Y = y$, the output of the IM construction is the pair of functions $(\mathsf{bel}, \mathsf{pl})_y$.

2.5 IM validity

The output of the IM construction is a belief function, but in what sense is this belief function meaningful for inference? One important property the belief function should have is that it does not tend to assign high beliefs to assertions about θ which are false. The validity property makes this precise.

Definition 2. An IM is called *valid* if, for any assertion A about θ, the corresponding belief function satisfies

$$\sup_{(\theta,\sigma):\theta\notin A} \mathsf{P}_{Y|(\theta,\sigma)}\{\mathsf{bel}_Y(A; \mathcal{S}) \geq 1-\alpha\} \leq \alpha, \quad \forall\, \alpha \in (0,1). \tag{8}$$

In other words, the IM is valid if, for any assertion A, $\mathsf{bel}_Y(A; \mathcal{S})$ is stochastically no larger than $\mathsf{Unif}(0,1)$ as a function of Y when A is false.

Martin and Liu (2013) showed that admissibility of $\mathcal{S}$ is sufficient for validity of the corresponding IM. The only downside is that there are many predictive random sets that give valid IMs. In that case, it is natural to look for a "best" one; see Definition 3.

3 General optimality results

3.1 Setup

Consider a general setup with association $X = a(\theta, U)$, where $X \in \mathbb{X}$ is the observable data, $\theta \in \Theta$ is the unknown parameter of interest, and U is an unobservable auxiliary variable, with distribution P_U defined on a space $\mathbb{U}$; precise measure-theoretical statement of our assumptions is given in the Supplementary Material. The corresponding sampling model is denoted by $\mathsf{P}_{X|\theta}$. We assume that the predictive random set is admissible, so that the IM's belief function bel_x is valid in the sense of Definition 2. Then, for a given assertion A about θ, an optimal predictive random set $\mathcal{S}$, if it exists, makes $\mathsf{bel}_X(A; \mathcal{S})$ stochastically as large as possible as a function of $X \sim \mathsf{P}_{X|\theta}$, for all $\theta \in A$ (Martin and Liu 2013). The following definition makes this formal.

Definition 3. For a given association and assertion A, if $\mathcal{S}$ and $\mathcal{S}'$ are two valid predictive random sets, then we say that $\mathcal{S}$ *is at least as efficient as* $\mathcal{S}'$ (with respect to A) if

$$\mathsf{bel}_X(A;\mathcal{S}) \geq_{\text{st}} \mathsf{bel}_X(A;\mathcal{S}'), \text{ as a function of } X \sim \mathsf{P}_{X|\theta},\ \theta \in A. \tag{9}$$

A predictive random set $\mathcal{S}$ is called optimal (with respect to A) if (9) holds for all valid predictive random sets $\mathcal{S}'$.

Given $\Theta_x(u) = \{\theta : x = a(\theta, u)\}$ from the A-step, define the *a-event* as

$$\mathbb{U}_A(x) = \{u \in \mathbb{U} : \Theta_x(u) \subseteq A\}. \tag{10}$$

This is just the set of all u that support the truthfulness of the assertion "$\theta \in A$" for a given x. We require that $\mathbb{U}_A(x)$ be appropriately measurable, and the precise conditions are given in the Supplementary Material. The a-events will be crucial in the construction of the focal elements of the optimal predictive random set.

As discussed in Section 1, towards developing an optimal predictive random set for a collection of assertions, simultaneously, we consider breaking these assertions down into their basic building blocks. Starting from the most basic building block, a simple assertion, we characterize the optimal predictive random sets and show how these can be combined to create optimal predictive random sets for more general kinds of assertions. Since the distribution of an admissible predictive random set is determined by (5), our focus throughout is on the construction of the focal elements.

3.2 Simple assertions

An assertion A is *simple* if the collection of a-events $\mathbb{U}_A \equiv \{\mathbb{U}_A(x) : x \in \mathbb{X}\}$ is nested, i.e., for any pair $x, x' \in \mathbb{X}$, we have either $\mathbb{U}_A(x) \subseteq \mathbb{U}_A(x')$ or $\mathbb{U}_A(x) \supseteq \mathbb{U}_A(x')$; otherwise, the assertion is called *complex*. The following results shows that, if A is simple, then the optimal predictive random set (relative to A), as in Definition 3, readily obtains by taking the support $\mathbb{S} = \mathbb{U}_A$ and the distribution $\mathsf{P}_{\mathcal{S}}$ to satisfy (5).

Theorem 1. *For a simple assertion A, the optimal predictive random set $\mathcal{S}$ with respect to A is supported on the nested collection $\mathbb{U}_A$ with distribution satisfying* (5).

Proof. See Appendix A.1. □

As an example, consider $X \sim \mathsf{N}(\theta, 1)$, with association $X = \theta + U$, $U \sim \mathsf{N}(0, 1)$. The assertion $A = \{\theta > 0\}$ is simple. To see this, write the a-event:

$$\mathbb{U}_A(x) = \{u : x = \theta + u \text{ for some } \theta > 0\} = (-\infty, x).$$

Then, clearly, if $x < x'$, then $\mathbb{U}_A(x) \subset \mathbb{U}_A(x')$; so, the a-events are nested. The belief function based on the optimal predictive random set $\mathcal{S}$ in Theorem 1 is

$$\mathsf{bel}_x(A;\mathcal{S}) = \mathsf{P}_{\mathcal{S}}\{\mathcal{S} \subset \mathbb{U}_x(A)\} = \mathsf{P}_U\{\mathbb{U}_A(x)\} = \Phi(x),$$

where Φ is the standard normal distribution function. Given the optimality result, it is not a coincidence that $\mathsf{pl}_x(A^c;\mathcal{S}) = 1 - \Phi(x)$ is the p-value for the uniformly most powerful test of $H_0 : \theta \leq 0$ versus $H_1 : \theta > 0$; see Martin and Liu (2014).

Unfortunately, simple assertions are insufficient for practical applications. For example, in the normal mean example above, we might be interested in the assertion $A = \{\theta \neq 0\}$ or, more generally, we might be interested in several assertions simultaneously, each of which might be simple or not. The regression application features all of these cases. (More generally, even if there is only one assertion A, IM efficiency considerations require that one think about A and A^c simultaneously, so an understanding of how to handle multiple assertions is fundamental.) The next two sections discuss how to extend the basic optimality results to these important more general cases.

3.3 Complex assertions

To motivate our developments, reconsider the normal mean example above. This time, suppose we are interested in the assertion $A = \{\theta \neq 0\}$. This is a union of two disjoint simple assertions, namely, $A_1 = \{\theta < 0\}$ and $A_2 = \{\theta > 0\}$. The optimal predictive random set with respect to A_1 is efficient for A_1 but very inefficient for A_2; likewise, for the optimal predictive random set with respect to A_2. Since we are interested in $A_1 \cup A_2$ we require a predictive random set that is as efficient as possible for both A_1 and A_2. For the general case of a complex assertion written as a union of two simple assertions, an intuitively reasonable strategy is to consider a predictive random set whose focal elements are intersections of the focal elements of the optimal predictive random sets with respect to the two individual simple assertions. The following result says that the optimal predictive random set for the complex assertion must be of this form. (The results below extend to more than two simple component assertions, but two-component assertions are general enough for our purposes here.)

Theorem 2. *Let $A = A_1 \cup A_2$ be a complex assertion, where A_1 and A_2 are simple assertions. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be the optimal predictive random sets for A_1 and A_2, respectively, as in Theorem 1. For any predictive random set $\mathcal{T}$, there exists an $\mathcal{S}$, whose focal elements are intersections of the focal elements of $\mathcal{S}_1$ and $\mathcal{S}_2$, such that $\mathcal{S}$ is at least as efficient as $\mathcal{T}$ with respect to A.*

Proof. See Appendix A.2. □

This result simplifies the search for optimal predictive random sets with respect to a complex assertion. It does not completely resolve the problem, however, since there are many choices of predictive random sets with intersection focal elements. In the normal mean problem, for example, the focal elements for the optimal predictive random sets with respect to A_1 and A_2 are one-sided intervals. Therefore, the focal elements of the optimal predictive random set for $A = A_1 \cup A_2$ must be nested intervals, but should the intervals be symmetric or asymmetric? See Section 3.5.

3.4 Multiple complex assertions

Suppose there are multiple assertions, $\{A_j : j \in J\}$, to be considered simultaneously, where $A_j = A_{j1} \cup A_{j2}$, $j \in J$, decomposes as a union of two disjoint simple assertions. Each simple component has an optimal predictive random set according to Theorem 1, and the corresponding optimal predictive random set, $\mathcal{S}_j$, for the union A_j has intersection

focal elements according to Theorem 2. As before, intuition suggests that the optimal predictive random set for $\{A_j : j \in J\}$ would be supported on intersections of the focal elements for the individually optimal $\mathcal{S}_j$, $j \in J$. To justify this intuition, we need a way to measure the efficiency of a predictive random set in this multiple-assertion context.

Definition 4. An assertion A is said to be *generated by* $\{A_j : j \in J\}$ if A can be written as a union of some or all of the A_js. Then a predictive random set $\mathcal{S}$ is at least as efficient as $\mathcal{S}'$ with respect to $\{A_j : j \in J\}$ if (9) holds for all A generated by $\{A_j : j \in J\}$.

The following result, a generalization of Theorem 2, shows that a restriction to predictive random sets supported on intersection focal elements is without loss of efficiency.

Theorem 3. *Let $\{A_j : j \in J\}$ be a collection of assertions, where $A_j = A_{j1} \cup A_{j2}$ partitions as a union of disjoint simple assertions. Let $\mathcal{S}_j$ be the optimal predictive random set for A_j, $j \in J$. Then, for any predictive random set $\mathcal{T}$, there exists an $\mathcal{S}$, whose focal elements are intersections of the focal elements of the $\mathcal{S}_j$s, such that $\mathcal{S}$ is at least as efficient as $\mathcal{T}$ with respect to $\{A_j : j \in J\}$.*

Proof. See Appendix A.3. □

As in the previous section, this result simplifies the search for an optimal predictive random set but does not completely resolve the problem, since it does not determine the particular shape of the focal elements. Therefore, we need to push further.

3.5 Balance and optimality

To resolve the ambiguity about the shape of the optimal focal elements, here we will introduce some special structure to the problem. Suppose that there exists transformations of X that do not fundamentally change the inference problem. As a simple example, if $X \sim \mathsf{N}(\theta, 1)$ and the association of interest is $A = \{\theta \neq 0\}$, then it is clear that changing the sign of X should not have an impact on how much support there is for A. In other words, this normal mean problem is invariant to sign changes. More generally, let $\mathscr{G}$ be a group of bijections g from $\mathbb{X}$ to itself; as is customary, we shall write gx for the image of x under g, rather than $g(x)$. Practically, the transformations in $\mathscr{G}$ represent symmetries of the inference problem, i.e., nothing fundamental about the problem changes if gX is observed instead of X. The key technical assumption here is that each g commutes with the association mapping a, i.e.,

$$g\, a(\theta, u) = a(g\theta, gu), \quad \forall\ (\theta, u), \quad \forall\ g \in \mathscr{G}. \tag{11}$$

Here we are implicitly assuming that $\mathscr{G}$ also acts upon the parameter space Θ and the auxiliary variable space $\mathbb{U}$. This can be relaxed by introducing groups acting on Θ and $\mathbb{U}$, respectively, different from (but homomorphic to) $\mathscr{G}$, but we will not need this extra notational complexity for our application here. The above condition is different from that which defines the usual group transformation models in that, in the right-hand side, g also acts on u. In fact, our variable selection application will not fit the usual group transformation structure, unless the design is orthogonal.

We shall also require that the assertions respect these symmetries. Let A be an assertion generated the collection $\{A_j : j \in J\}$ of complex assertions. Consider the

subgroup of $\mathscr{G}$ to which A is invariant, and write $\mathscr{G}_A = \{g \in \mathscr{G} : gA = A\}$. The intuition is that the inference problem for θ, at least as it concerns the assertion A, is unchanged if the problem is transformed by $g \in \mathscr{G}_A$.

Before proceeding, it may help to see a simple example. Consider, again, the normal mean problem, with $X = \theta + U$, $U \sim \mathsf{N}(0,1)$, and assertion $A = \{\theta \neq 0\}$. Then changing the sign of X will not affect the problem. So, we can take $\mathscr{G}$ to consist of the identity mapping and $x \mapsto -x$, and clearly (11) holds; also, $\mathscr{G}_A = \mathscr{G}$.

Moving on, recall the a-event $\mathbb{U}_A(x) = \{u\colon\ x = a(\theta, u) \text{ for some } \theta \in A\}$. In this case, an equivariance property follows immediately from the commutativity property (11):

$$\mathbb{U}_A(gx) = g\mathbb{U}_A(x), \quad \forall\, g \in \mathscr{G}_A.$$

That is, given A, transforming $x \to gx$, for $g \in \mathscr{G}_A$, and then solving for u is equivalent to solving for u with the given x and then transforming $u \to gu$.

We now have the necessary structure to help specify an optimal predictive random set. It suffices to focus on predictive random sets which are admissible. So, let $\mathcal{S}$ be admissible and, for simplicity, suppose we can write the collection of closed nested focal elements as $\mathbb{S} = \{S_r : r \in [0,1]\}$, where r corresponds to the set's P_U-probability; in particular, $\mathsf{P}_U(S_r) = 1 - r$. Then we have the following useful representation.

Lemma 1. $\mathsf{P}_{X|\theta}\{\mathsf{bel}_X(A;\mathcal{S}) > 1 - r\} = \mathsf{P}_{X|\theta}\{\mathbb{U}_A(X) \supset S_r\}$.

Proof. See Appendix A.4. □

From Lemma 1 and the equivariance property above, if $g \in \mathscr{G}_A$, then

$$\mathsf{P}_{X|\theta}\{\mathsf{bel}_{gX}(A;\mathcal{S}) > 1 - r\} \equiv \mathsf{P}_{X|\theta}\{\mathbb{U}_A(gX) \supset S_r\} = \mathsf{P}_{X|\theta}\{\mathbb{U}_A(X) \supset g^{-1}S_r\}.$$

Since the understanding is that the inference problem, at least as it concerns the assertion A, is unchanged by transformations $X \to gX$ for $g \in \mathscr{G}_A$, it is reasonable to require that the distribution of $\mathsf{bel}_X(A;\mathcal{S})$ be invariant to $\mathscr{G}_A$. The previous display reveals that the way to achieve belief function invariance is to require the focal elements of the predictive random set to be invariant to $\mathscr{G}$. This leads to the following notion of *balance*.

Definition 5. The predictive random set $\mathcal{S}$ is said to be *balanced with respect to* A if each focal element $S \in \mathbb{S}$ satisfies $gS = S$ for all $g \in \mathscr{G}_A$. Moreover, $\mathcal{S}$ is said to be *balanced* if the aforementioned invariance holds for all $g \in \mathscr{G}$.

Balance itself is a reasonable property, given that the transformations are, by definition, irrelevant to the inference problem. It is also interesting and practically beneficial that balance can be checked without doing any probability calculations; however, the focal elements $\mathbb{S}$ and the transformations $\mathscr{G}$ depend on the model.

Beyond the intuitive appeal of balance, we claim that balance leads to a particular form of optimality. Recall that, for IM efficiency, the general goal is to choose the predictive random set to make the belief function stochastically large when the assertion is true. With this in mind, we propose the following notion of *maximin optimality*.

Definition 6. A predictive random set $\mathcal{S}^\star$, with focal elements $\{S_r^\star : r \in [0,1]\}$, is *maximin optimal* if it maximizes

$$\min_{g \in \mathscr{G}_A} \mathsf{P}_{X|\theta}\{\mathsf{bel}_{gX}(A;\mathcal{S}) > 1 - r\} \equiv \min_{g \in \mathscr{G}_A} \mathsf{P}_{X|\theta}\{\mathbb{U}_A(X) \supset g^{-1}S_r\}$$

over all admissible $\mathcal{S}$ with focal elements $\{S_r : r \in [0,1]\}$, uniformly for all A generated by $\{A_j : j \in J\}$, for all $\theta \in A$, and for all $r \in [0,1]$.

Theorem 4. *If a predictive random set $\mathcal{S}$ is balanced in the sense of Definition 5, then it is maximin optimal in the sense of Definition 6.*

The main idea in the proof, presented in Appendix A.5, is the notion of the "core" S° of a given focal element S. In particular, the core is defined as $S^\circ = \bigcap_{g \in \mathscr{G}_A} gS$, and the proof relies on the fact that it is both balanced and contained in each gS.

4 Model uncertainty in regression

4.1 Optimal IM for assertions about the model

Recall the dimension-reduced association $\hat{\theta} = \theta + \hat{\sigma}U$ in (4), where U is a p-vector distributed as $\mathsf{P}_U = \mathsf{t}_p(0, L, n-p-1)$, and L is a matrix with ones on the diagonal. The intercept has been marginalized out, so each of $\theta_1, \ldots, \theta_p$ can be considered zero or non-zero; the goal is to summarize the uncertainty about these possible models.

Consider the collection of complex assertions $A_j = \{\theta : \theta_j \neq 0\}$, $j = 1, \ldots, p$. Consider first a particular A_j. This can be written as $A_j = A_{j1} \cup A_{j2}$, where $A_{j1} = \{\theta : \theta_j < 0\}$ and $A_{j2} = \{\theta : \theta_j > 0\}$. We claim that these sub-assertions are both simple in the sense of Section 3.2. Take A_{j1}, for example. Then the corresponding a-event is

$$\mathbb{U}_{A_{j1}}(y) = \{u : \hat{\theta} - \hat{\sigma}u \in A_{j1}\} = \{u : u_j > T_j\},$$

where $T_j = \hat{\theta}_j / \hat{\sigma}$. This a-event is nested because it shrinks and expands monotonically as a function of T_j. The same is true for A_{j2} and for all the other $j = 1, \ldots, p$. Therefore, by Theorem 1, the optimal predictive random sets for the individual sub-assertions A_{j1} and A_{j2} are each supported on collections of half hyper-planes. Next, it follows from Theorem 2 that the optimal predictive random set $\mathcal{S}_j$ for the complex assertion A_j is supported on intersections of half hyper-planes, i.e., cylinders $\{u : a_j \leq u_j \leq b_j\}$. If we are considering $\{A_j : j = 1, \ldots, p\}$ simultaneously, then it follows from Theorem 3 that the optimal predictive random set is supported on boxes—intersections of p marginal cylinders—in $\mathbb{R}^p$. The remaining question is what shape should the boxes be.

Towards optimality, we need to consider what transformations leave the model uncertainty problem invariant in the sense of Section 3.5. As in the simple normal mean example discussed previously, sign changes to coordinates of $\hat{\beta}$ are irrelevant. In addition, the labeling of the variables $j = 1, \ldots, p$ is arbitrary, so permutations of the variable labels are also irrelevant. This suggests we consider the group $\mathscr{G}$ of *signed permutations*; that is, each $g \in \mathscr{G}$ acts on a p-vector x by matrix multiplication (on the left), where the matrix factors as a product of a diagonal matrix with ± 1 on the diagonal and a permutation matrix. It is clear that the association commutes with $\mathscr{G}$ in the sense of (11). With the group $\mathscr{G}$ specified, it is also clear what shape of boxes the optimal predictive random set should be supported on. According to Definition 5, a balanced predictive random set should have focal elements—in this case, shaped like boxes in $\mathbb{R}^p$—that are invariant to the transformations in $\mathscr{G}$. The only such boxes are hyper-cubes centered at the origin.

Corollary 1. *The admissible hyper-cube predictive random set $\mathcal{S}$, given by*

$$\mathcal{S} = \{u \in \mathbb{R}^p : \|u\|_\infty \leq \|U\|_\infty\}, \quad U \sim \mathsf{P}_U = \mathsf{t}_p(0, L; n-p-1), \tag{12}$$

is balanced in the sense of Definition 5 and maximin optimal for $\{A_j : j = 1, \ldots, p\}$ in the sense of Definition 6.

4.2 IM output and its interpretation

Now that we have specified the optimal predictive random set for the P-step, the C-step proceeds just as before. It is helpful, however, to describe the IM output and its interpretation as it pertains to model uncertainty quantification. To start, we need a bit more notation. Let $\mathcal{J} \subseteq \{1, 2, \ldots, p\}$ denote the collection of indices corresponding to the truly non-zero coefficients, i.e., $\theta_j \neq 0$ for all $j \in \mathcal{J}$. Also, for a generic subset $J \subseteq \{1, 2, \ldots, p\}$, let θ_J denote the corresponding sub-vector of θ.

Consider the following assertion about the model:

$$C_J = \{\mathcal{J} \subseteq J\} = \{\theta : \theta_{J^c} = 0\}.$$

Note that $C_J^c = \bigcup_{j \in J^c} A_j$, so this new assertion is generated by the model assertions $A_1, \ldots, A_p$ from before. Using the optimal predictive random set $\mathcal{S}$ in (12), we have

$$\mathsf{pl}_Y(C_J; \mathcal{S}) = \mathsf{P}_U\{\|U\|_\infty > \|T_{J^c}\|_\infty\} = 1 - F(\|T_{J^c}\|_\infty), \tag{13}$$

where F is the distribution function of $\|U\|_\infty$ when $U \sim \mathsf{t}_p(0, L; n-p-1)$. The distribution function F can be evaluated via Monte Carlo or the method of Genz and Bretz (2002), implemented in the R package `pmvt`.

Toy Example. For the illustrative example in Section 1.2, where $p = 2$, the error variance $\sigma^2 = 1$ was known, and the design matrix was identity, the distribution function F has a simpler form, namely, $F(z) = \{1 - 2\Phi(-z)\}^2$, $z \geq 0$. Then the plausibility functions computed in that example, and summarized in Figure 1, are based on the expression (13) with this distribution function and T as the pair of observations (X_1, X_2).

How to interpret the plausibility (13) above? On one hand, it is clear both from intuition and from the formula that $\mathsf{pl}_Y(C_J; \mathcal{S})$ is non-decreasing in J and that the full model has plausibility 1. On the other hand, since the optimal hyper-cube predictive random set has positive volume, the belief in any proper sub-model is zero. So, to borrow the terminology from Dempster (2008), larger models have more "don't know" probability, which is consistent with idea that larger models have more parameters to estimate and are more difficult to interpret. Consequently, only the relatively small plausibilities, which correspond to relatively small models, can be interpreted. To help understand what "relatively small" means, we have the following calibration result.

Theorem 5. *If C_J is true, then $\mathsf{pl}_Y(C_J; \mathcal{S})$ is stochastically no smaller than $\mathsf{Unif}(0, 1)$.*

Proof. It is clear that, C_J is true, i.e., if $\theta_{J^c} = 0$, then $\|U\|_\infty$ is stochastically no smaller than $\|T_{J^c}\|_\infty$. Then the result follows from the fact that $F(\|U\|_\infty) \sim \mathsf{Unif}(0, 1)$. □

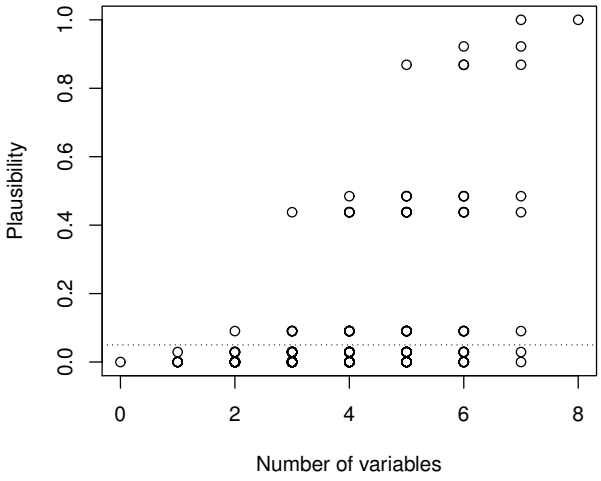

Figure 2: Plot of the model plausibility values, arranged by model size, for the prostate cancer data example studied in Section 4.2.

The result in Theorem 5 explains why the IM plausibility's distribution function falls on or below the diagonal line in Panels (a) and (c) in Figure 1. Moreover, optimality of the predictive random set suggests that the IM should also be efficient, i.e., if C_J is false, then $\mathsf{pl}_Y(C_J; \mathcal{S})$ should be stochastically smaller than $\mathsf{Unif}(0,1)$, which is the conclusion in Panels (b) and (d) of the same figure.

For further illustration, consider the prostate cancer data analyzed by Tibshirani (1996) and others. This study examined the association between the prostate specific antigen (PSA) level and some clinical measures among men who were about to receive a radical prostatectomy. There were $n = 97$ men with $p = 8$ predictors, four of which, including cancer volume (`lcavol`), prostate weight (`lweight`), capsular penetration (`lcp`), and benign prostatic hyperplasia amount (`lbph`), were log transformed. The other four predictors were age, seminal vestical invasion (`svi`), Gleason score (`gleason`), and percentage Gleason scores 4 or 5 (`pgg45`). The response variable is the log PSA level.

We can compute the plausibilities (13) for each J, and the results are summarized in Figure 2, where the plausibilities are arranged by the corresponding model size. That is, for each $s = 0, 1, \ldots, 8$, there are $\binom{8}{s}$ plausibility values displayed vertically. As mentioned previously, only the relatively small plausibilities in this figure can be interpreted and, since we are interested in models which are both small and sufficiently plausible, we focus on maximum plausibility for each s. The maximum plausibilities for $s = 1$ and $s = 2$ are both rather small; the maximum value (0.09) in the $s = 2$ column corresponds to the model that contains variables `lcavol` and `svi` only. If we move to $s = 3$ variables, then the largest plausibility value jumps to 0.44, corresponding to the model that adds the variable `lweight` to the previous two-variable model. In fact, this latter three-variable model is the one that is selected by lasso. The plot reveals that there is a lot more uncertainty, i.e., "don't know," for model assertions that allow for the possibility that `lweight` is included in the true model. We conclude that data only provides evidence to support the claim that the true model is a subset of $\{$`lcavol`, `svi`$\}$; see Section 4.4.

4.3 On post-selection inference

Using the balanced hyper-cube predictive random set $\mathcal{S}$ in Corollary 1, we can construct a plausibility function for singletons $\{\theta\}$:

$$\mathsf{pl}_Y(\theta;\mathcal{S}) := \mathsf{pl}_Y(\{\theta\};\mathcal{S}) = 1 - F\big(\|\hat{\sigma}^{-1}(\hat{\theta}-\theta)\|_\infty\big).$$

This plausibility function satisfies $\mathsf{pl}_Y(\hat{\theta};\mathcal{S}) = 1$ and is decreasing in θ away from $\hat{\theta}$ with hyper-cube shaped contours. Consequently, by thresholding the plausibility function at level $\alpha \in (0,1)$, we obtain the plausibility region

$$\{\theta : \mathsf{pl}_Y(\theta;\mathcal{S}) \geq \alpha\} = \big\{\theta : F\big(\|\hat{\sigma}^{-1}(\hat{\theta}-\theta)\|_\infty\big) \leq 1-\alpha\big\}.$$

Since the predictive random set $\mathcal{S}$ is admissible, it follows from the general IM theory that the plausibility region above has nominal frequentist coverage $1-\alpha$. Moreover, the shape of these plausibility regions is a hyper-cube, with side lengths characterized by quantiles of the ℓ_∞-norm of multivariate Student-t random vectors.

An important consequence of our focus on validity simultaneously across a collection of assertions is that a naive projection of these plausibility regions to any sub-model $J \subseteq \{1,2,\ldots,p\}$ gives a new hyper-cube plausibility region that also has nominal frequentist coverage. Since this conclusion holds uniformly in J, it also holds if J is chosen based on data. Such considerations are relevant in the context of post-selection inference. As Berk et al. (2013) argue, when data is used to select a model, then one cannot use the model-specific distribution theory for valid inference. This leads to the fundamental question of how to achieve valid inference when the model is chosen based on data. They propose a procedure they call "POSI" for valid post-selection inference, and it turns out that their procedure is identical to that obtained by projection of the above plausibility region to a sub-model selected by data. The take-away message here is that, by insisting on simultaneous validity over all relevant assertions, the IM approach can automatically handle the challenging post-selection inference problem. This is a consequence of our focus on the main goal of valid uncertainty quantification about the model.

4.4 IM-driven variable selection

Section 4.1 showed how to optimally quantify model uncertainty, and the previous section explained how this implies valid post-selection inference. It is also possible to develop an IM-driven variable selection procedure having desirable properties. We want to be clear that valid uncertainty quantification about the model is our primary goal, and the result of a variable selection procedure is only one kind of summary of the IM output.

Intuitively, a model or collection of variables J is supported by data Y if $\mathsf{pl}_Y(C_J;\mathcal{S})$ is not too small. In other words, good models are those which are "sufficiently plausible" given data. Following this idea, and the remarks in Section 4.2, an IM-driven approach for variable selection would be to fix $\alpha \in (0,1)$ and then pick the smallest collection J such that $\mathsf{pl}_Y(C_J;\mathcal{S}) > \alpha$. That is, define

$$\hat{J}_\alpha(Y) = \text{smallest set } J \text{ such that } \mathsf{pl}_Y(C_J;\mathcal{S}) > \alpha. \tag{14}$$

We claim that the IM-driven procedure (14) satisfies a *selection validity* property:

$$\mathsf{P}_{Y|\mathcal{J}}\{\hat{J}_\alpha(Y) \subseteq \mathcal{J}\} \geq 1-\alpha, \quad \forall\, \mathcal{J} \subseteq \{1,2,\ldots,p\}. \tag{15}$$

Theorem 6. *The procedure* (14) *has the selection validity property* (15).

Proof. $\mathsf{pl}_Y(C_{\mathcal{J}}; \mathcal{S}) > \alpha$ implies $\hat{J}_\alpha(Y) \subseteq \mathcal{J}$, so apply Theorem 5. □

To implement the above procedure, it is not necessary to evaluate the plausibility function at C_J for each $J \subseteq \{1, 2, \ldots, p\}$. In fact, the plausibility depends on the value of $\|T_{J^c}\|_\infty$, so we only need to look at p different models, based on a sorting of the t-statistics by their magnitude. Let π be a permutation that ranks the T values in descending order according to their magnitudes, i.e., $|T_{\pi(1)}| > |T_{\pi(2)}| > \cdots > |T_{\pi(p)}|$, and then compute

$$\mathsf{mpl}_Y(s) = \max_{J:|J|=s} \mathsf{pl}_Y(C_J; \mathcal{S}) = 1 - F(|T_{\pi(s+1)}|), \quad s = 0, 1, \ldots, p-1, \tag{16}$$

and $\mathsf{mpl}_Y(p) \equiv 1$. This is a naive "marginal" plausibility function for the *size* of the model (Shafer 1987, Sec. G). If $s_\alpha^\star$ is the smallest s such that $\mathsf{mpl}_Y(s) > \alpha$, then

$$\hat{J}_\alpha(Y) = \pi^{-1}(\{1, \ldots, s_\alpha^\star\}). \tag{17}$$

For example, for the prostate cancer data discussed in Section 4.2, based on the summary in Figure 2, if set the cutoff at $\alpha = 0.05$, then the IM-based rule (14) selects the two-variable model that includes `lcavol` and `svi`.

A simulation study is performed to assess the performance of this variable selection procedure. We generate 1000 data sets consisting of observations $Y \sim \mathsf{N}_n(X\beta, \sigma^2 I_n)$, where the rows of the predictor variable matrix are draws from $\mathsf{N}_p(0, \Omega)$ with an autoregressive correlation structure, i.e., $\Omega_{jk} = \rho^{|j-k|}$ for all j, k. This model was used by Tibshirani (1996) in his simulation study. We take $\sigma = 3$ and $\rho = 0.5$ and consider two different scenarios:

Scenario 1. $p = 7$, $\beta = (3, 1.5, 0, 0, 2, 0, 0)^\top$;
Scenario 2. $p = 20$, $\beta = (0.85\, 1_{10}^\top, 0_{10}^\top)^\top$.

The relative comparisons based on simulations under some different settings were the same as here so these are not presented here. For each scenario, we vary the sample size n from 50 to 5000. For each simulated data set, the IM-driven variable selection procedure is applied. Performance of this approach is compared to the lasso, where the tuning parameter is chosen using the 10-fold cross validation, the adaptive lasso, AIC, and BIC. AIC and BIC variable selection is performed using exhaustive search of all possible models. Results of these approaches are summarized in Figures 3 and 4 for the two scenarios. Specifically, we plot the percentage of selecting the true model, where all important variables are selected (Panel A), a subset of the true model, where some but not all important variables and no unimportant variables are selected (Panel B), a superset of the true model, where all important variables and at least one unimportant variable are selected (Panel D), and the other models, where at least one important variable is not selected and at least one unimportant variable is selected (Panel E). In order to show the selection validity of the IM-driven approach, we also plot the percentage of selecting the true model or a subset of the true model, where some or all important variables are selected while no unimportant variables are selected (Panel C).

In both scenarios, the IM-driven approach performs well in terms of how frequently it selects the true model—similar to BIC but better than lasso and AIC. When it misses

the true model, the IM-driven approach tends to select a subset of the important variables in both scenarios. Other approaches, on the other hand, are more likely to select unimportant variables. This is especially the case in Scenario 2 with a larger number of variables, where these approaches miss important variables and incorrectly select unimportant variables (Panel E). Overall, the IM-based approach maintains the 95% level for selecting the important variables without picking up unimportant variables (Panel C), a consequence of Theorem 6. The other approaches, except for BIC and adaptive lasso, are generally far off the mark. That BIC and adaptive lasso can reach above 95% for selecting true model or subset of true model is due to their asymptotic variable selection consistency property, not because they are properly calibrated at any fixed n.

The take-away message is that the variable selection procedure (14) derived from our IM is surely competitive with some of the standard methods used in the literature. Other ad hoc selection procedures can be used which may beat our IM-driven variable selection procedure with respect to some criterion, but these procedures will always fall short of providing uncertainty quantification about the model.

5 Discussion

This paper introduces the concept of multiple simultaneous assertions in the IM framework, and develops a theory of optimal predictive random sets. These general principles are then applied, in a regression setting, to provide a valid quantification of the uncertainty in the model. To our knowledge, the IM approach is the only method known to satisfy this important property. The IM's simultaneous validity property sheds light on the post-selection inference problem, and also leads to a variable selection procedure with desirable frequentist error rate control. Our simulation study demonstrates that, overall, the proposed IM-based procedure performs as good or better than several standard methods, which suggests that the IM's emphasis on valid probabilistic uncertainty quantification does not come at the cost of decreased efficiency. Extending the developments here for application in generalized linear model settings is a focus of ongoing research.

The IM approach has already been applied to other problems that involve multiplicity, such as large-scale multinomial inference in genome wide association studies and multiple testing problems (Liu and Xie 2014a,b). We expect that the optimality considerations here, applied to those problems, would lead to some overall improvements. Perhaps the most important question is how to extend the developments in this paper to the high-dimensional regression context. The basic principles of optimal predictive random sets developed here for simultaneous complex assertions are not specific to the $p < n$ case; however, the initial dimension reduction steps do not carry over directly to the $p \gg n$ case. We expect that once the A-step can be completed for $p \gg n$, ideas similar to those presented here can be applied.

Acknowledgements

The authors thank the Editor, Associate Editor, and referees for their valuable comments on an earlier draft of this manuscript. This work was partially supported by the U.S. National Science Foundation: DMS–1007678, DMS–1208833, and DMS–1208841.

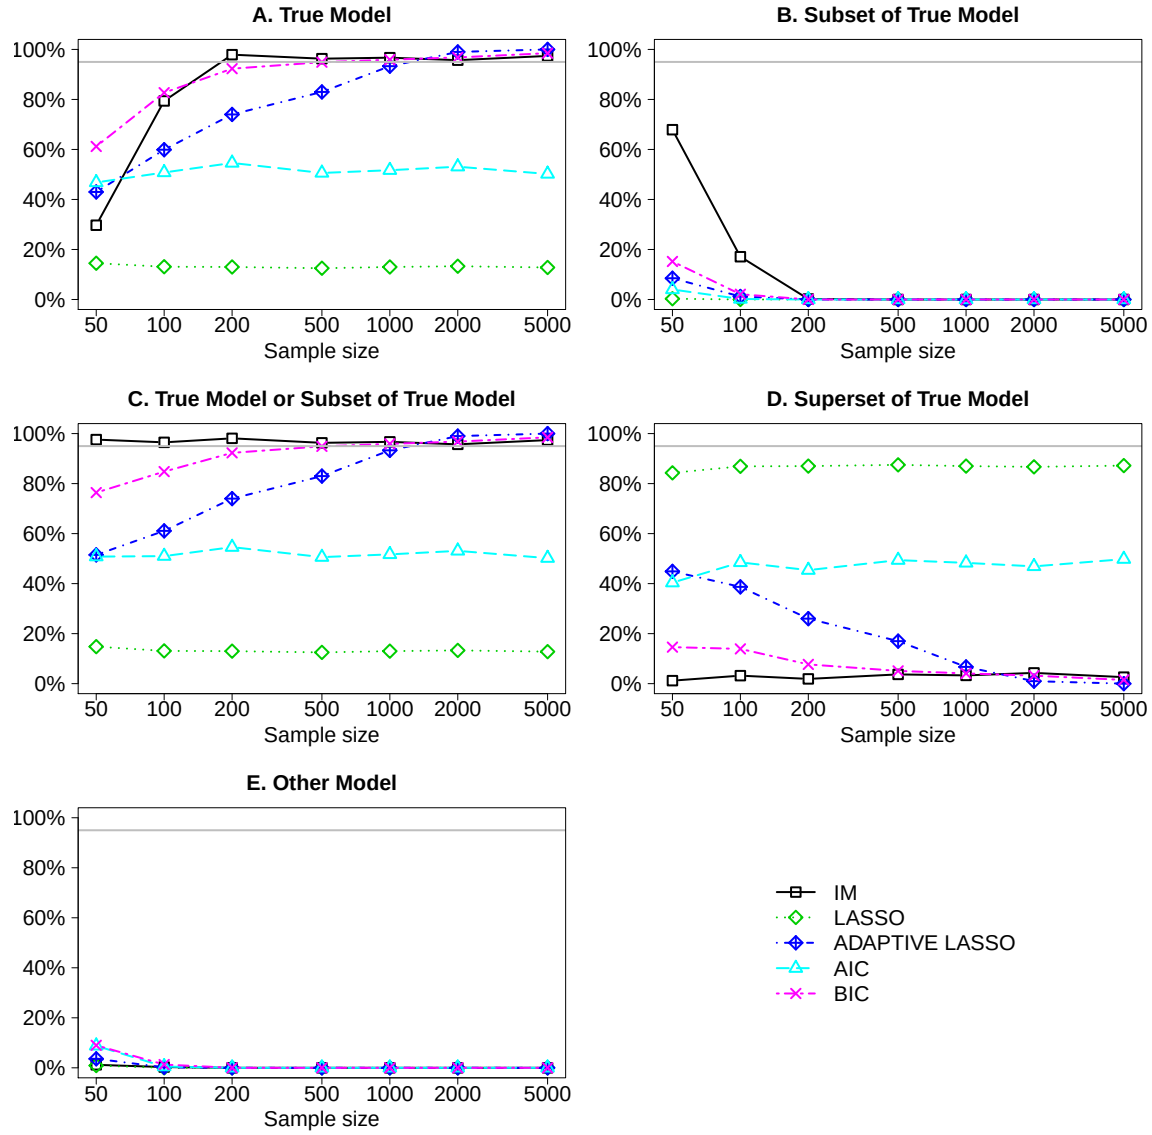

Figure 3: Percentage of selecting the true model (panel A), subset of the true model where some but not all important variables is selected (panel B), true model or subset of true model where some or all important variables are selected (panel C), superset if true model where all important variables and at least one unimportant variable are selected (panel D), and other model where at least one important variable is not selected and at least one unimportant variable is selected (panel E) for Scenario 1.

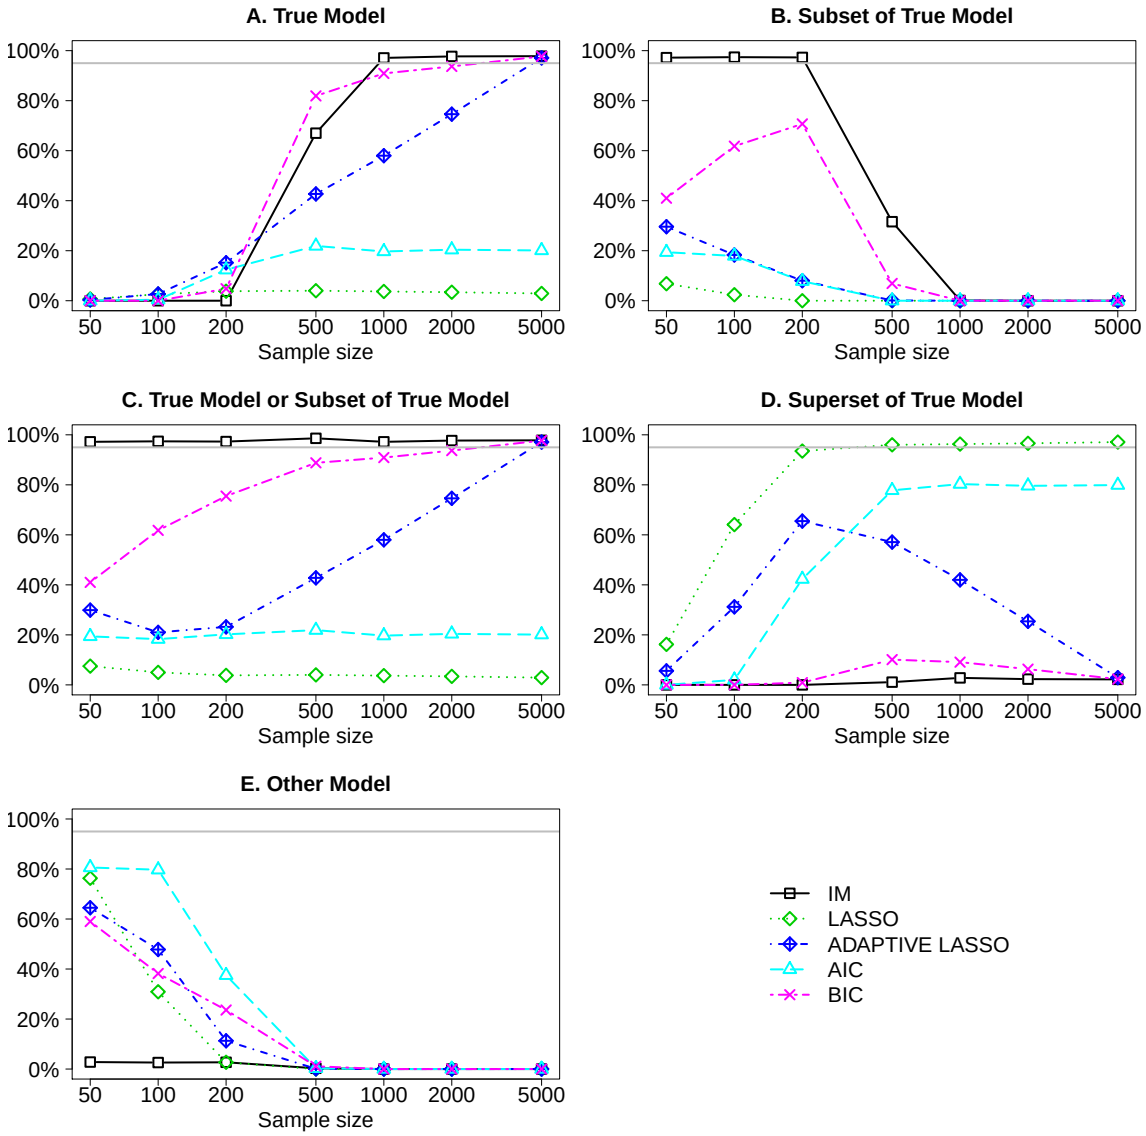

Figure 4: Same explanation as for Figure 3 but for Scenario 2.

The authors also thank Professor Jan Hannig for suggesting a possible connection between our work and that of Berk et al. (2013).

A Proofs

A.1 Proof of Theorem 1

Proposition 1 in Martin and Liu (2013) shows that $\mathsf{bel}_x(A;\mathcal{S}) \leq \mathsf{P}_U\{\mathbb{U}_A(x)\}$, for any admissible $\mathcal{S}$ and any x. By assumption, the collection $\{\mathbb{U}_A(x) : x \in \mathbb{X}\}$ is nested, so it can be taken as the support of an admissible predictive random set. In this case, since $\mathsf{P}_\mathcal{S}$ satisfies (5), we have that $\mathsf{bel}_x(A;\mathcal{S}) = \mathsf{P}_U\{\mathbb{U}_A(x)\}$ for all x. Therefore, the belief function attains its upper bound for each x, hence, it is optimal.

A.2 Proof of Theorem 2

Since A_1 and A_2 are simple assertions, the respective optimal predictive random sets $\mathcal{S}_1$ and $\mathcal{S}_2$ have focal elements given by $\mathbb{U}_{A_1}(x)$ and $\mathbb{U}_{A_2}(x)$, as x ranges over $\mathbb{X}$. Without loss of generality, assume that the predictive random set $\mathcal{T}$ for $A = A_1 \cup A_2$ is admissible. That is, assume $\mathcal{T}$ has a nested support $\mathbb{T}$ and that $\mathsf{P}_\mathcal{T}$ satisfies (5). Define

$$S_j(T) = \text{closure}\Big\{ \bigcap_{y:\mathbb{U}_{A_j}(y)\supset T} \mathbb{U}_{A_j}(y) \Big\} \quad j = 1, 2, \quad T \in \mathbb{T}.$$

Collect intersections of these sets, $\mathbb{S} = \{S_1(T) \cap S_2(T) : T \in \mathbb{T}\}$. Since $\mathbb{T}$ is nested and the function $T \mapsto S_1(T) \cap S_2(T)$ is monotone, the collection $\mathbb{S}$ is also nested. Define a new predictive random set $\mathcal{S}$, supported on $\mathbb{S}$, with the natural measure $\mathsf{P}_\mathcal{S}$ as in (5). This predictive random set satisfies the conditions stated in the theorem, i.e., admissible and has focal elements as intersections of the respective optimal predictive random set focal elements. We need to show that $T \subset \mathbb{U}_A(x)$ if and only if $S(T) \subset \mathbb{U}_A(x)$, where $S(T) = S_1(T) \cap S_2(T)$. One direction is easy, since it is clear that $T \subset S(T)$. For the other direction, we want to show that $T \subset \mathbb{U}_A(x)$ implies $S(T) \subset \mathbb{U}_A(x)$. Since $\mathbb{U}_{A_1}(x)$ and $\mathbb{U}_{A_2}(x)$ are disjoint, and the union is $\mathbb{U}_A(x)$, $T \subset \mathbb{U}_A(x)$ implies that either $T \subset \mathbb{U}_{A_1}(x)$ or $T \subset \mathbb{U}_{A_2}(x)$. By definition of $S_1(T)$ and $S_2(T)$, it follows that

$$S_1(T) \subset \mathbb{U}_{A_1}(x) \quad \text{or} \quad S_2(T) \subset \mathbb{U}_{A_2}(x).$$

In either case, $S(T) = S_1(T) \cap S_2(T)$ is contained in $\mathbb{U}_A(x) = \mathbb{U}_{A_1}(x) \cup \mathbb{U}_{A_2}(x)$, which completes the argument that $\{T \subset \mathbb{U}_A(x)\}$ and $\{S(T) \subset \mathbb{U}_A(x)\}$ are equivalent. Finally, since $T \subset S(T)$, we get $\mathsf{P}_U(T) \leq \mathsf{P}_U(S(T))$ for all $T \in \mathbb{T}$ and, therefore,

$$\mathsf{P}_\mathcal{T}\{\mathcal{T} \subset \mathbb{U}_A(x)\} = \sup_{T:T\subset\mathbb{U}_A(x)} \mathsf{P}_U(T) \leq \sup_{T:S(T)\subset\mathbb{U}_A(x)} \mathsf{P}_U(S(T)) = \mathsf{P}_\mathcal{S}\{\mathcal{S} \subset \mathbb{U}_A(x)\}.$$

The left-hand side is $\mathsf{bel}_x(A;\mathcal{T})$ and the right-hand side is $\mathsf{bel}_x(A;\mathcal{S})$, and the inequality holds for all x, completing the proof.

A.3 Proof of Theorem 3

The proof here is similar to that of Theorem 2. Consider collection $\{A_j : j \in J\}$ of complex assertions, where each A_j can be written as a union of disjoint simple assertions, i.e., $A_j = A_{j1} \cup A_{j2}$, where $A_{j1} \cap A_{j2} = \varnothing$ and the a-events $\mathbb{U}_{A_{j1}}(\cdot)$ and $\mathbb{U}_{A_{j2}}(\cdot)$ are nested. By Theorem 2, we know that the optimal predictive random sets for A_j, $j \in J$, have focal elements $\mathbb{S}_j = \{S_j(v) : v \in V\}$, indexed by a set V, which are intersections of a-events. That is, $S_j(v)$ is (the closure of) $\mathbb{U}_{A_{j1}}(x_{j1}(v)) \cap \mathbb{U}_{A_{j2}}(x_{j2}(v))$ for some $x_{j1}(v)$ and $x_{j2}(v)$.

Without loss of generality, assume that the candidate predictive random set $\mathcal{T}$ for the assertion A generated by $\{A_j : j \in J\}$ is admissible in the sense of Definition 1. Given a focal element $T \in \mathbb{T}$ of $\mathcal{T}$, define

$$\tilde{S}_j(T) = \bigcap_{v : S_j(v) \supset T} S_j(v), \quad j \in J.$$

Next, set $\tilde{S}(T) = \bigcap_{j \in J} \tilde{S}_j(T)$ and define $\mathbb{S} = \{\tilde{S}(T) : T \in \mathbb{T}\}$. Now take $\mathcal{S}$ to have support $\mathbb{S}$ and natural measure $\mathsf{P}_{\mathcal{S}}$ as in (5); this $\mathcal{S}$ satisfies the conditions of the theorem. It remains to show that $\mathcal{S}$ is more efficient than $\mathcal{T}$.

As in the proof of Theorem 2, we need to show that $T \subset \mathbb{U}_A(x)$ if and only if $\tilde{S}(T) \subset \mathbb{U}_A(x)$ for each x and $T \in \mathbb{T}$. By construction, $T \subset \tilde{S}(T)$, so one direction is handled. For the other direction, assume that $T \subset \mathbb{U}_A(x)$. Then $T \subset \mathbb{U}_{A_j}(x)$ for some $j \in J$ and, since A_j splits as a disjoint union of simple assertions, we get that $T \subset \mathbb{U}_{A_{j1}}(x)$ or $T \subset \mathbb{U}_{A_{j2}}(x)$. Then $\tilde{S}_j(T) \subset \mathbb{U}_{A_{j1}} \cap \mathbb{U}_{A_{j2}}(x)$ and, consequently, the no-bigger $\tilde{S}(T)$ must be a subset of $\mathbb{U}_A(x)$. Since $T \subset \mathbb{U}_A(x)$ if and only if $\tilde{S}(T) \subset \mathbb{U}_A(x)$, the claimed superiority of $\mathcal{S}$ to $\mathcal{T}$ follows just like in the last part of the proof of Theorem 2.

A.4 Proof of Lemma 1

Recall that $\mathsf{bel}_x(A; \mathcal{S})$ is defined as $\mathsf{P}_{\mathcal{S}}\{\mathcal{S} \subset \mathbb{U}_A(x)\}$. Since $\mathcal{S}$ has the natural measure (5), we can write, for any $b \in [0, 1]$,

$$\begin{aligned}
\mathsf{bel}_x(A; \mathcal{S}) > b &\iff \mathsf{P}_{\mathcal{S}}\{\mathcal{S} \subset \mathbb{U}_A(x)\} > b \\
&\iff \sup_{r : S_r \subset \mathbb{U}_A(x)} \mathsf{P}_U(S_r) > b \\
&\iff \sup_{r : S_r \subset \mathbb{U}_A(x)} (1 - r) > b \\
&\iff \mathbb{U}_A(x) \supset S_{1-b}.
\end{aligned}$$

Therefore, we have that

$$\mathsf{P}_{X|\theta}\{\mathsf{bel}_X(A; \mathcal{S}) > b\} = \mathsf{P}_{X|\theta}\{\mathbb{U}_A(X) \supset S_{1-b}\}, \quad \forall\, b \in [0, 1],$$

which proves the claim, with $r = 1 - b$.

A.5 Proof of Theorem 4

Take any predictive random set $\mathcal{S}$ as in Theorem 3, and let S be a generic focal element. Take any assertion A generated by $\{A_j : j \in J\}$. Define the core of S as $S^\circ = \bigcap_{g \in \mathscr{G}_A} gS$;

note that S° is balanced and satisfies $S^\circ \subset gS$ for all $g \in \mathscr{G}_A$. Then,

$$\mathsf{P}_{X|\theta}\{\mathbb{U}_A(X) \supset S^\circ\} \geq \mathsf{P}_{X|\theta}\{\mathbb{U}_A(X) \supset gS\}, \quad \forall\, g \in \mathscr{G}_A, \quad \forall\, \theta \in A,$$

with strict inequality in general. Maximin optimality requires that we choose the focal elements to maximize the minimum (over g) of the right-hand side of the above display. However, we can clearly attain the upper bound above by taking the focal element S equal to its core, i.e., balanced. Therefore, balance implies maximin optimality.

References

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory (Tsahkadsor, 1971)*, pages 267–281. Akadémiai Kiadó, Budapest.

Berger, J. O., Bernardo, J. M., and Sun, D. (2009). The formal definition of reference priors. *Ann. Statist.*, 37(2):905–938.

Berk, R., Brown, L., Buja, A., Zhang, K., and Zhao, L. (2013). Valid post-selection inference. *Ann. Statist.*, 41(2):802–837.

Bühlmann, P. (2013). Statistical significance in high-dimensional linear models. *Bernoulli*, 19(4):1212–1242.

Clyde, M. and George, E. I. (2004). Model uncertainty. *Statist. Sci.*, 19(1):81–94.

Dempster, A. P. (2008). The Dempster–Shafer calculus for statisticians. *Internat. J. Approx. Reason.*, 48(2):365–377.

Ermini Leaf, D. and Liu, C. (2012). Inference about constrained parameters using the elastic belief method. *Internat. J. Approx. Reason.*, 53(5):709–727.

Fraser, D. A. S. (2011). Is Bayes posterior just quick and dirty confidence? *Statist. Sci.*, 26(3):299–316.

Fraser, D. A. S., Reid, N., Marras, E., and Yi, G. Y. (2010). Default priors for Bayesian and frequentist inference. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(5):631–654.

Genz, A. and Bretz, F. (2002). Comparison of methods for the computation of multivariate t probabilities. *J. Comput. Graph. Statist.*, 11(4):950–971.

Hannig, J., Iyer, H., Lai, R. C. S., and Lee, T. C. M. (2015). Generalized fiducial inference: A review. *J. Amer. Statist. Assoc.*, to appear.

Lai, R. C. S., Hannig, J., and Lee, T. C. M. (2015). Generalized fiducial inference for ultrahigh dimensional regression. *J. Amer. Statist. Assoc.*, 110:760–772.

Liu, C. and Xie, J. (2014a). Large scale two sample multinomial inferences and its applications in genome-wide association studies. *Internat. J. Approx. Reason.*, 55(1, part 3):330–340.

Liu, C. and Xie, J. (2014b). Probabilistic inference for multiple testing. *Internat. J. Approx. Reason.*, 55(2):654–665.

Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. (2014). A significance test for the lasso. *Ann. Statist.*, 42(2):413–468.

Martin, R. and Liu, C. (2013). Inferential models: A framework for prior-free posterior probabilistic inference. *J. Amer. Statist. Assoc.*, 108(501):301–313.

Martin, R. and Liu, C. (2014). A note on p-values interpreted as plausibilities. *Statist. Sinica*, 24:1703–1716.

Martin, R. and Liu, C. (2015a). Conditional inferential models: Combining information for prior-free probabilistic inference. *J. R. Stat. Soc. Ser. B*, 77(1):195–217.

Martin, R. and Liu, C. (2015b). Marginal inferential models: Prior-free probabilistic inference on interest parameters. *J. Amer. Statist. Assoc.*, 110:1621–1631.

Martin, R., Mess, R., and Walker, S. G. (2015). Empirical Bayes posterior concentration in sparse high-dimensional linear models. *Bernoulli*, to appear; `arXiv:1406.7718`.

Molchanov, I. (2005). *Theory of Random Sets*. Probability and Its Applications (New York). Springer-Verlag London Ltd., London.

O'Hara, R. B. and Sillanpää, M. J. (2009). A review of Bayesian variable selection methods: what, how and which. *Bayesian Anal.*, 4(1):85–117.

Schwarz, G. (1978). Estimating the dimension of a model. *Ann. Statist.*, 6(2):461–464.

Shafer, G. (1987). Belief functions and possibility measures. In Bezdek, J. C., editor, *The Analysis of Fuzzy Information, Vol. 1: Mathematics and Logic*, pages 51–84. CRC.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B*, 58(1):267–288.

Tibshirani, R. (2011). Regression shrinkage and selection via the lasso: a retrospective. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 73(3):273–282.

Xie, M. and Singh, K. (2013). Confidence distribution, the frequentist distribution of a parameter – a review. *Int. Statist. Rev.*, 81(1):3–39.

Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In *Bayesian Inference and Decision Techniques*, volume 6 of *Stud. Bayesian Econometrics Statist.*, pages 233–243. North-Holland, Amsterdam.

Zhang, Z., Xu, H., Martin, R., and Liu, C. (2011). Inferential models for linear regression. *Pak. J. Statist. Oper. Res.*, 7:413–432.

Zou, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 67(2):301–320.

Supplementary Material

S Measure-theoretic details

S.1 Random sets and the admissibility condition

Molchanov (2005) gives a comprehensive treatment of the theory of random sets. Our goal here is present the minimal amount of technical details necessary to understand our analysis involving predictive random sets. To start, the standard theory of random sets focuses on the case of closed random sets, i.e., random sets whose values are closed sets (with probability 1). Write $\mathbb{S}$ for the support of our random set $\mathcal{S}$, a collection of subsets of the space $\mathbb{U}$; we assume that $\mathbb{S}$ contains both $\varnothing$ and $\mathbb{U}$. Let $\mathbb{U}$ be a separable metric space and take $\mathscr{U}$ to be a σ-algebra of subsets that contains all the closed subsets of $\mathbb{U}$. Assume that each set, or focal element, $S \in \mathbb{S}$ is closed, relative to the topology on $\mathbb{U}$.

To define the random set $\mathcal{S}$, consider a probability space $(\Omega, \mathscr{A}, \mathsf{P})$, and let $\mathcal{S}$ be mapping from Ω to $\mathbb{S}$ which is measurable in the sense that

$$\{\omega : \mathcal{S}(\omega) \cap K \neq \varnothing\} \in \mathscr{A}, \quad \text{for all compact } K \subseteq \mathbb{U}.$$

We define the distribution $\mathsf{P}_{\mathcal{S}}$ of $\mathcal{S}$ as the push-forward measure $\mathsf{P}\mathcal{S}^{-1}$. Also, for compact K, the event $\{\mathcal{S} \subset K\}$ is measurable, so the probability $\mathsf{P}_{\mathcal{S}}\{\mathcal{S} \subset K\}$ in (5) and also below makes sense. Moreover, since $\mathbb{U}$ is separable and $\mathcal{S}$ is closed, the indicator stochastic process $\{I_{\mathcal{S}}(u) : u \in \mathbb{U}\}$ is separable and, by Proposition 4.10 in Molchanov (2005), the distribution $\mathsf{P}_{\mathcal{S}}$ is determined by probabilities assigned to the events $\{\mathcal{S} \subset K\}$, and can be extended to include arbitrary sets K.

An important class of examples are the predictive random sets described in Corollary 1 of Martin and Liu (2013); the "default" predictive random set that has been frequently used in the IM literature, as well as the optimal predictive random set employed in this paper, are members of this class. In particular, start by taking the probability space $(\Omega, \mathscr{A}, \mathsf{P})$ to be $(\mathbb{U}, \mathscr{U}, \mathsf{P}_U)$. Next, for a continuous function $h : \mathbb{U} \to \mathbb{R}$, define the set-valued mapping ψ_h on $\mathbb{U}$ as

$$\psi_h(u) = \{u' \in \mathbb{U} : h(u') \leq h(u)\}, \quad u \in \mathbb{U}.$$

Then the focal elements $S \in \mathbb{S}$ are closed level sets of the function h, and if $\mathscr{U}$ is rich enough to contain all the closed sets, then ψ_h is a measurable function and, consequently, $\mathcal{S} = \psi_h(U)$, for $U \sim \mathsf{P}_U$, is a closed random set.

In the IM context, as a consequence of the results in Martin and Liu (2013), we focus only on predictive random sets with nested support $\mathbb{S}$, where nested means that, for any two focal elements, one is a subset of the other. Such supports can, and usually are, constructed as in the small example above, but requiring that the function h acts like a metric, i.e., h has a unique minimizer and that $h(u)$ is increasing as u moves further away from that minimizer. The default predictive random set, for the case $\mathbb{U} = [0, 1]$, takes $h(u) = |u - 0.5|$, which is of the form just described. Using the terminology in Shafer (1987), we could call a predictive random set with nested support *consonant* and, in addition to the admissibility properties demonstrated in Martin and Liu (2013), consonant random sets have the simplest distributional properties.

In our restriction to admissible predictive random sets, as described in Definition 1, there is a more pressing concern, namely, does there exist a predictive random set $\mathcal{S}$, supported on a given nested collection $\mathbb{S}$ of closed P_U-measurable subsets, such that

$$\mathsf{P}_{\mathcal{S}}\{\mathcal{S} \subset K\} = \sup_{S \in \mathcal{S}: S \subset K} \mathsf{P}_U(S).$$

Specifically, while the right-hand side above is a well-defined quantity, it is not obvious that there is a random set $\mathcal{S}$ with distribution $\mathsf{P}_{\mathcal{S}}$ that satisfies the above equality. However, the famous Choquet capacity theorem (Molchanov 2005, Theorem 1.13) can be applied to show existence of such an $\mathcal{S}$. In particular, from the above display, one can identify the condition that the corresponding capacity function of $\mathcal{S}$ must satisfy and, from the close connection with the probability measure P_U, the required upper semi-continuity and completely alternating properties can be checked relatively easily.

S.2 Measurability assumptions

In addition to measurability of set-valued mappings, as discussed above, there are questions about some more familiar forms of measurability, pertaining to the a-events $\mathbb{U}_A(x)$ for a given assertion $A \subseteq \Theta$. Recall that the auxiliary variable space $\mathbb{U}$ is equipped with a σ-algebra $\mathscr{U}$ of P_U-measurable subsets, assumed to contain all the closed sets. Next, equip the sample space $\mathbb{X}$ with a σ-algebra $\mathscr{X}$ and a σ-finite measure μ, where $\mathsf{P}_{X|\theta} \ll \mu$ for each θ. Then the two key measurability assumptions are:

- $\mathbb{U}_A(x) \in \mathscr{U}$ for all relevant A and for μ-almost all x, and
- $\{x : \mathbb{U}_A(x) \supset S\} \in \mathscr{X}$ for all focal elements $S \in \mathbb{S}$.

General sufficient conditions can be given based properties of the association mapping $x = a(\theta, u)$ and of the assertion A. It is relatively easy to check the above conditions directly in a particular example and, moreover, the conditions might fail only in non-standard problems.